arXiv:1603.01419v2 [cs.RO] 30 Jun 2016

Methods for Stochastic Collection and Replenishment (SCAR) optimisation for persistent autonomy

Andrew W. Palmer, Andrew J. Hill, Steven J. Scheding*

Australian Centre for Field Robotics, The University of Sydney, NSW, Australia

Abstract

Consideration of resources such as fuel, battery charge, and storage space, is a crucial requirement for the successful persistent operation of autonomous systems. The Stochastic Collection and Replenishment (SCAR) scenario is motivated by mining and agricultural scenarios where a dedicated replenishment agent transports a resource between a centralised replenishment point to agents using the resource in the field. The agents in the field typically operate within fixed areas (for example, benches in mining applications, and fields or orchards in agricultural scenarios), and the motion of the replenishment agent may be restricted by a road network. Existing research has typically approached the problem of scheduling the actions of the dedicated replenishment agent from a short-term and deterministic angle. This paper introduces a method of incorporating uncertainty in the schedule optimisation through a novel prediction framework, and a branch and bound optimisation method which uses the prediction framework to minimise the downtime of the agents. The prediction framework makes use of several Gaussian approximations to quickly calculate the risk-weighted cost of a schedule. The anytime nature of the branch and bound method is exploited within an MPC-like framework to outperform existing optimisation methods while providing reasonable calculation times in large scenarios.

Keywords: Persistent autonomy, scheduling, uncertainty, refuelling, recharging

*Email address: {a.palmer;a.hill;s.scheding}@acfr.usyd.edu.au

1. Introduction

Achieving persistent autonomy requires careful management of finite resources such as fuel, battery charge, and storage space. While there are some scenarios where agents are able to continue operating by collecting energy from the environment, such as gliding Unmanned Aerial Vehicles (UAVs) exploiting thermals to remain airborne [1], in the majority of cases the agent uses a resource which must be replenished to enable persistence. Early work on replenishment involved the robot returning to a charging dock to facilitate autonomous recharging [2, 3], which now forms the basis of recharging in some commercial systems such as the iRobot Roomba [4] and InTouch Health RP-7i [5]. However, there are many cases in which the use of a dedicated replenishment agent can provide direct benefit. Examples include refuelling or recharging of UAVs in flight, refuelling of satellites in orbit, and data collection from underwater wireless sensor networks. Using a dedicated agent to replenish the resource of each agent in the field enables the agents to be smaller, simpler, and cheaper [6].

The Stochastic Collection and Replenishment (SCAR) problem was first introduced by the authors in [7] as a generic resource management scenario that is motivated by scenarios commonly found in mining and agricultural environments. Possible resources include fuel, battery charge, food, and water for the replenishment case, and electronic data, mined ore, and harvested fruit for the collection case. In these scenarios, a dedicated replenishment agent transports the resource between a centralised replenishment point and the agents operating in the field. The motion of the replenishment agent is typically restricted by roads or other obstacles, and the agents in the field operate within fixed areas such as benches in mining scenarios, and fields and orchards in agricultural scenarios. The SCAR scenario differs from similar research in several key ways. Firstly, the examined literature assumes that each agent is visited only once. When considering persistent scenarios, neglecting future visits may lead to sub-optimal decision making. Secondly, the replenishment agent is generally assumed to have sufficient capacity such that it will not run out of the resource. Again, this is not the case for persistent autonomy. Finally and most importantly, the examined literature assumes that the problem is deterministic. In practice, this is not the case as agent parameters such as speed and resource usage rate will have elements of uncertainty.

A prediction framework for SCAR scenarios was developed by the authors

in [7] which was able to quickly calculate the risk-weighted cost of a schedule for the replenishment agent. This was then used within an A* optimisation method in [8] to schedule the actions of the replenishment agent in small SCAR scenarios. This paper introduces an improved version of the prediction framework and develops a branch and bound optimisation method for use in large SCAR scenarios where optimisation time is limited. The specific contributions of this paper include:

- a novel framework for predicting the risk-weighted cost of a schedule;
- a novel Gaussian approximation to the inverse of a Gaussian distributed random variables;
- a novel Gaussian approximation to the generalised rectified Gaussian distribution;
- an anytime branch and bound optimisation method; and
- a computational comparison of the proposed approaches with the state of the art.

The prediction framework is accurate and computationally efficient. This is significant as complex objective functions that incorporate uncertainty have, thus far, required the use of Monte Carlo simulation to calculate an accurate estimate of the expected value. Combined with the branch and bound approach, the proposed methods are shown to outperform the existing approaches in terms of minimising the downtime of the agents in the field.

The rest of this paper is structured as follows: Section 2 discusses the related literature, and a description of the SCAR scenario is introduced in Section 3. The improved prediction framework is outlined in Section 4, and approximations for using Gaussian distributed random variables within the prediction framework are developed in Section 5. The optimisation methods are introduced in Section 6. The prediction and optimisation methods are then compared in a computational study in Section 7, and conclusions and future work are presented in Section 8.

2. Related literature

The main replenishment and collection scenarios in the literature are refuelling or aircraft and satellites, recharging of ground and aerial robots, and

collecting data from wireless sensor networks. In general, these are framed as NP-hard combinatorial optimisation problems that resemble classical problems such as the restricted Travelling Salesman Problem (TSP) with time windows, or the parallel machine manufacturing job shop problem.

The aerial refuelling problem was examined in [9–12]. Jin et al. [9] used a recursive dynamic programming approach to determine the optimal order for refuelling the aircraft to minimise the priority weighted time of refuelling for each aircraft. Such an approach is not applicable to SCAR scenarios as the optimal future decision in a SCAR scenario is dependent on the sequence of decisions leading to it. Kaplan and Rabadi [10, 11] developed heuristic and meta-heuristic approaches. In [10] they found that the meta-heuristic was outperformed by the heuristic in large scenarios, and in [11] they used the heuristic to generate an initial schedule for the meta-heuristic to improve performance.

An assumption of the work of Jin et al. and Kaplan and Rabadi was that the aircraft being refuelled were in formation behind the tanker aircraft, meaning that the time between refuelling each aircraft was determined purely by the set-up time required by that particular aircraft. Barnes et al. [12] considered the inter-theatre aerial refuelling, where the flight time of the tanker aircraft between aircraft to be refuelled was treated as sequence-dependent set-up times. This is similar to how the travel between satellites was treated for refuelling a constellation of satellites [13]. The solution methods used in these cases were heavily tailored towards the specific scenario under consideration and are not generalisable to SCAR scenarios.

Recharging of ground and aerial robots was examined in [14–16]. Kannan et al. [14] developed a market-based solution for determining a good recharging strategy, but only examined scenarios with a single user agent, while Litus et al. [15] focussed on determining rendezvous locations given an optimal recharging order. Mathew et al. [16] developed a receding horizon approach for simultaneously calculating a recharging order and rendezvous locations. These papers assumed that the recharging agent had sufficient capacity to fully recharge all of the user agents without itself having to return to a charging point, ignoring a critical aspect for persistent operation.

The literature on collection scenarios has focused on the use of data mules to collect data from wireless sensor networks. This strategy can be particularly useful for underwater wireless sensor networks, used for long-term monitoring of coral reefs, where wireless communications over long distances can be difficult. Vasilescu et al. [17] used an Autonomous Underwater Vehicle

(AUV) to travel to each sensor node and retrieve the collected data. Similar scenarios are presented in [18–24]. The problem is generally treated as a variant of the TSP with the aim of minimising the total distance travelled, the data latency, or the total schedule time.

Most of the above literature does not approach this problem from a persistent autonomy perspective—the optimisation methods generally consider visiting each agent or node only once, assume that the replenishment agent has sufficient or unlimited capacity, and ignore time-varying effects such as variable replenishment times and deadlines. In addition, all of the above literature bar [16] ignore uncertainty. In [16], arbitrary safety margins were used to reduce plan failure, with no consideration of the actual risk associated with each task.

Uncertainty in scheduling has received limited study, even in the classical manufacturing scheduling literature, due to its difficulty in comparison to deterministic problems [25]. Typical approaches to incorporating uncertainty include chance constraints [26], Monte Carlo simulation [27], using conservative estimates of the uncertain parameters [28], replanning frequently [29], or simply ignoring the uncertainty [30]. Of these approaches, only chance constraints and Monte Carlo simulation allow risk to be incorporated into the optimisation.

The prediction framework presented in the authors' previous work [7] produced a similar result to Monte Carlo simulation. Instead of sampling the probability distributions, however, it used analytical and approximation methods to propagate the entire probability distribution, resulting in a calculation time that was orders of magnitude faster than the Monte Carlo approach. This was used in [8] within an A* optimisation approach to evaluate each schedule under consideration. Using the prediction framework to incorporate risk was shown to outperform an A* approach that ignored the risk. While it was sufficient for small SCAR scenarios, this approach is infeasible for larger scenarios due to the size of the search space.

Optimisation methods which are generally used in larger scenarios include heuristics, meta-heuristics, and anytime combinatorial optimisation methods. The Apparent Tardiness Cost (ATC) heuristic and simulated annealing meta-heuristic were used by Kaplan and Rabadi [10] for the aerial refuelling problem. They found that the ATC heuristic outperformed simulated annealing in larger scenarios, and in later work combined the two methods to find better solutions [11]. Anytime combinatorial optimisation methods include branch and bound, and anytime A* methods such as ARA*. ARA* uses

inadmissible heuristics to quickly find sub-optimal solutions, before refining the solution by moving towards an admissible heuristic [31], while branch and bound prunes parts of the solution tree which have an estimated lower cost bound that is higher than the cost of the current solution [32]. Due to the difficulty in deriving heuristics that can be used in ARA*, branch and bound has been chosen as the solution method in this paper.

3. The SCAR scenario

SCAR scenarios are motivated by mining and agricultural scenarios such as replenishing excavators and drills using fuel and water trucks, and collecting harvested crops and picked fruit. These scenarios consist of multiple user agents, such as excavators and harvesters, that either consume or collect a resource over time. As they have a limited capacity of the resource, a dedicated replenishment agent, such as a fuel truck, rendezvous with the user agents and replenishes their supply of the resource. This enables the user agents to continue operating in the field without having to return to fixed replenishment infrastructure. Collection scenarios are identical to replenishment scenarios if the resource under consideration is storage space. An example SCAR scenario is shown in Figure 1.

The fleet of user agents is heterogeneous, and, for the scenarios under consideration, there is a centralised replenishment point such as a refuelling station or silo. The parameters of the user agents, replenishment agent, and replenishment point, such as their speed, set-up and pack-up times, and resource usage rates, are stochastic. The user agents operate in defined areas such as benches on a mine site for drills and excavators, and fields for tractors and harvesters in agricultural scenarios. The distances between the operational areas of the user agents are generally much larger than the size of those areas, and any variations in the travel times of the replenishment agents due to the movements of the user agents are assumed to be accounted by the uncertain travel speed, and set-up and pack-up times of the replenishment agent.

Note that user agents are not required to remain in their operational area indefinitely. User agents that move to new benches or fields within the scope of a schedule can be incorporated by treating the time that they leave their current location as the latest time that it can be replenished at that location. Similarly, the time that it arrives at its new location can be treated as the ready time for that agent. Given uncertainty in the arrival time of the

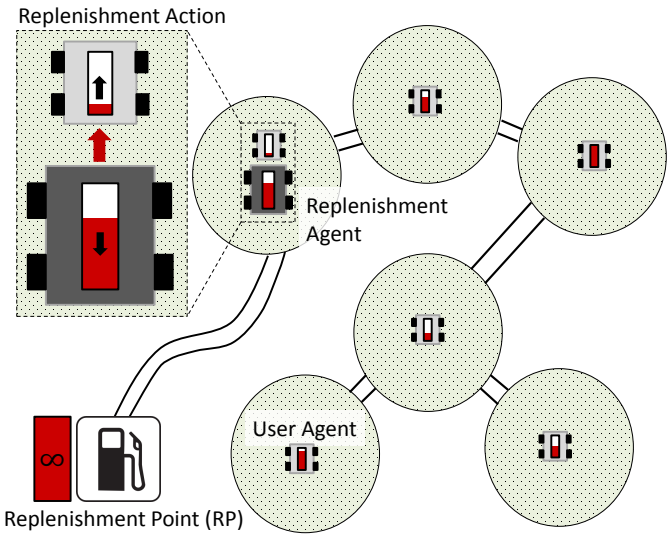

Figure 1: An example SCAR scenario. A replenishment agent (dark grey) travels to, and replenishes, the user agents (light grey) operating in the field. The replenishment agent must return to the replenishment point (bottom left) to replenish its supply of the resource. The replenishment point has infinite capacity of the resource. In this example, the travel of the replenishment agent is restricted by roads between the operational areas of the user agents. Inset: The replenishment agent transfers the resource to the user agent, diminishing its supply of the resource and increasing the resource reserves of the user agent.

replenishment agent, a probabilistic check can be used to evaluate whether the replenishment agent will arrive at the location before the deadline or after the ready time with sufficient probability for it to be allowed as a valid task. In this way, operational areas can be added or removed from the network. Locations along roads are not considered valid locations for the replenishment agent to service a user agent as blocking the road creates a hazardous situation.

It is assumed that the user agents can recover from exhausting their supply of the resource, i.e. they enter a safe zero-resource state. This is modelled in this paper as a soft deadline which has a cost that increases linearly with the time that the user agents are not operational. A hard deadline would correspond to, for example, a UAV exhausting its supply of fuel mid-flight, causing the loss of the agent.

The following subsections introduce the system parameters, variables, and constraints, and the optimisation objectives. Parameters and variables are treated as either known to a high degree of certainty (denoted by a lower case letter), or as random variables (denoted by an upper case letter). The uncertain parameters and variables in this paper are assumed to be Gaussian distributed random variables, and several approximations are introduced in Section 5 to facilitate the mathematical operations performed on them. Other probability distributions can be used provided that the appropriate operations and approximations exist. Units do not matter so long as they are consistent.

3.1. Parameters, variables, and constraints

For each user agent, i, in an n-user agent system:

- $c_{u,i}$ is the resource capacity
- $l_{u,i}$ is the current resource level, where $l_{u,i} \in [0, c_{u,i}]$
- $L_{u,i}$ is the estimated future resource level, where $L_{u,i} \in [0, c_{u,i}]$
- $R_{u,i}$ is the resource usage rate
- w_i is a user-defined weight or priority for the user agent

It is assumed that the user agents continue consuming the resource while being replenished. If the current resource level of the user agent reaches 0, then the user agent ceases operation and incurs downtime.

For the replenishment agent:

- c_a is the resource capacity
- l_a is the current resource level, where $l_a \in [0, c_a]$
- L_a is the estimated future resource level, where $L_a \in [0, c_a]$
- R_a is the resource replenishment rate into the user agent
- D_{sa} is the duration of time required for the replenishment agent to set-up at a user agent
- D_{pa} is the duration of time required for the replenishment agent to pack-up at a user agent
- V_a is the velocity

The replenishment agent is assumed to have a separate supply of fuel or battery charge for its own operation that is replenished in parallel at the replenishment point. It is also assumed to be able to service only one user agent at a time. To replenish a user agent, the replenishment agent must first travel to the user agent and set up before commencing the replenishment. After it has either fully replenished the user agent or exhausted its own supply of the resource, the replenishment agent must pack up before travelling to the next task.

Finally, for the replenishment point:

- R_r is the resource replenishment rate into the replenishment agent
- D_{sr} is the duration of time required for the replenishment agent to set-up at the replenishment point
- D_{pr} is the duration of time required for the replenishment agent to pack-up at the replenishment point

For the replenishment agent to be replenished, it must first travel to the replenishment point, then set up, be fully replenished by the replenishment point, and then pack up before moving onto the next task. It cannot service a user agent while being replenished by the replenishment point. Note that the set-up and pack-up times at the replenishment point are different to those at the user agents.

The distance between the replenishment agent and user agent i is denoted $s_{au,i}$, and the distance between the replenishment agent and the replenishment point is denoted s_{ar}. In the case where there are multiple routes between tasks, the replenishment agent is assumed to take the shortest (fastest) route.

3.2. Optimisation

The aim of the optimiser is to minimise the total downtime of they user agents by optimising the actions of the replenishment agent, where downtime is incurred when a user agent has exhausted its supply of the resource. More generally, the objective function is the total weighted downtime of the user agents, ζ:

$$\min \zeta = \sum_{i=1}^{n} w_i d_{c,i} \tag{1}$$

where $d_{c,i}$ is the total downtime incurred by user agent i. Given the uncertainty present in the agent parameters, the total downtime that is expected to be incurred for a given schedule can only be estimated. Two methods for estimating the downtime of the user agents for a given schedule of actions are given in the authors' previous work [7], and an improved method is developed in Section 4.

The total weighted downtime objective is only appropriate for planning over an infinite horizon. As has been demonstrated in the authors' previous work [8], the use of a ratio based objective function can produce better results when using combinatorial optimisation methods to optimise the actions of the replenishment agent. The ratio objective function, λ, is defined as:

$$\min \lambda = \frac{\zeta}{n d_m} \tag{2}$$

where d_m is the total time for the schedule to be executed. Provided the weights sum to n, then the objective function is bound between 0 and 1:

$$\text{if } \sum_{i=0}^{n-1} w_i = n \text{ then } 0 \leq \lambda \leq 1 \tag{3}$$

The ratio objective function enables comparison between schedules that have the same number of tasks but take different lengths of time to execute. A schedule consists of an ordered list of tasks for the replenishment agent to perform. The task of visiting the replenishment point is denoted by 0, while

the task of replenishing a user agent is denoted by its index number. An example schedule, $\boldsymbol{\theta}$, of a 4-user agent SCAR scenario is:

$$\boldsymbol{\theta} = (1, 0, 4, 2, 1, 4)$$

In this schedule, the first task, $\boldsymbol{\theta}_1$, involves the replenishment agent replenishing the first user agent before visiting the replenishment point for task $\boldsymbol{\theta}_2$. For tasks $\boldsymbol{\theta}_3, \boldsymbol{\theta}_4, \boldsymbol{\theta}_5$ and $\boldsymbol{\theta}_6$, it would replenish agents 4, 2, 1 and 4 again in that order. Note that user agents 1 and 4 appear multiple times, and user agent 3 does not appear at all. Unlike existing replenishment scenarios, the assumptions of the SCAR scenario allow for user agents to be visited multiple times, or not at all within a given schedule.

The optimisation is performed within a framework similar to Model Predictive Control (MPC)—the optimiser returns a new task or schedule after each task is completed. Thus, unexpected changes to the system state are incorporated into the optimisation each time a task is performed.

4. Prediction framework

This section develops an improved version of the analytical prediction framework presented in [7]. The framework takes a schedule for the replenishment agent as input, and its aim is to predict the future resource levels of the user and replenishment agents, and to estimate the total weighted downtime of the user agents and the total time taken to execute the schedule. This framework will be used within the combinatorial optimisation methods developed in Section 6 to evaluate the objective function.

In [7], two continuous-time frameworks were developed—the first used a Monte Carlo approach to estimate the downtime of the user agents given the uncertainty in the system parameters, and the second used analytical and approximation methods to propagate the uncertainty instead of the sampling used in the Monte Carlo approach. The framework presented in Algorithm 1 builds on the second approach from [7]. It is agnostic to the type of probability distribution used, and Section 5 presents approximations that enable the use of Gaussian distributed random variables. To use other probability distributions with this framework, these approximations would need to be formulated for the cases where the appropriate operations do not exist.

On line 1, the total weighted downtime is initialised to 0 and the start time of the schedule is initialised to the current time. Note that the start time, T_l,

Algorithm 1: Prediction framework

`Analytical(`$\boldsymbol{\psi}_{initial}, \boldsymbol{\theta}$`)`

input : Current system state, $\boldsymbol{\psi}_{initial}$; schedule, $\boldsymbol{\theta}$; current time, t_{cur}

output: Ratio objective function, λ; state, $\boldsymbol{\psi}$

$\zeta \leftarrow 0$, $T_l \leftarrow t_{cur}$ `// initialise weighted downtime and start time`

while *there are tasks remaining in the schedule,* $\boldsymbol{\theta}$ **do**

  if $\boldsymbol{\theta}_1 = 0$ **then**

    $T_a \leftarrow T_l + \frac{s_{ar}}{V_a}$ `// calculate arrival time`

    $T_f \leftarrow T_a + D_{sr} + \frac{c_a - L_a}{R_r} + D_{pr}$ `// replenishment finish time`

    $L_a \leftarrow c_a$ `// reset resource level`

  else

    $i \leftarrow \boldsymbol{\theta}_1$ `// index of target user agent`

    $T_a \leftarrow T_l + \frac{s_{au,i}}{V_a}$ `// calculate arrival time`

    $T_{b,i} \leftarrow T_a + D_{sa}$ `// time after set-up`

    $T_{d,i} \leftarrow T_{f,i} + \frac{L_{u,i}}{R_{u,i}}$ `// deadline`

    $D_{c,i} \leftarrow T_{b,i} - T_{d,i}$ `// downtime`

    $E(\text{downtime}) \leftarrow \int_0^\infty d\ p(D_{c,i})\ \mathrm{d}d$ `// expected downtime`

    $\zeta \leftarrow \zeta + w_i E(\text{downtime})$ `// weighted downtime`

    $L_{u,i} \leftarrow (L_{u,i} - (T_{b,i} - T_{f,i})\, R_{u,i})^{\#}$ `// level before replenishment`

    $Q_{u,i} \leftarrow (c_{u,i} - L_{u,i})\, \frac{R_a}{R_a - R_{u,i}}$ `// replenishment quantity`

    $Q^*_{u,i} \leftarrow (Q_{u,i})^{*\leq L_a}$ `// adjusted quantity`

    $D_{r,i} \leftarrow \frac{Q^*_{u,i}}{R_a}$ `// replenishment duration`

    $L_{u,i} \leftarrow (L_{u,i} - (T_{b,i} - T_{f,i})\, R_{u,i})^{\#}$ `// user agent level`

    $L_a \leftarrow (L_a - Q_{u,i})^{\#}$ `// replenishment agent level`

    $T_{f,i} \leftarrow T_{b,i} + D_{r,i}$ `// update last replenishment time`

    $T_l \leftarrow T_{f,i} + D_{pa}$ `// time after pack-up`

  remove the task $\boldsymbol{\theta}_1$ from the schedule, $\boldsymbol{\theta}$

forall the *user agents,* i **do**

  $T_{d,i} \leftarrow T_{f,i} + \frac{L_{u,i}}{R_{u,i}}$ `// deadline`

  $D_{c,i} \leftarrow T_l - T_{d,i}$ `// downtime`

  $E(\text{downtime}) \leftarrow \int_0^\infty d\ p(D_{c,i})\ \mathrm{d}d$ `// expected downtime`

  $\zeta \leftarrow \zeta + w_i E(\text{downtime})$ `// weighted downtime`

$\boldsymbol{\psi} \leftarrow (L_a, L_{u,i}\ \forall i)$ `// predicted final state`

$\lambda \leftarrow \frac{\zeta}{nE(T_l - t_{cur})}$ `// ratio objective function`

is initialised as a random variable with no uncertainty. Then, while there are tasks remaining in the schedule, the first task in the schedule is evaluated. If the task is for the replenishment agent to visit the replenishment point (line 3), then the random variable describing the arrival time of the replenishment agent at the replenishment point, T_a, is calculated as per line 4. The time that the replenishment is completed at, T_f, is given on line 5, and the resource level of the replenishment agent is reset to its capacity on line 6.

If the task is instead to travel to and replenish a user agent, then the arrival time is calculated on line 9, and the time after the replenishment agent has set up, $T_{b,i}$, is calculated on line 10. The time at which the user agent would exhaust its supply of the resource if it were not replenished, $T_{d,i}$, is calculated on line 11, where $T_{f,i}$ is the time that the user agent was last replenished (see line 21). If the user agent has not been visited by a replenishment agent, then $T_{f,i} = t_{cur}$. The random variable describing the duration of downtime incurred by the user agent between when it was last replenished and the start time of the current replenishment action, $D_{c,i}$, is calculated on line 12.

The random variable $D_{c,i}$ is described by a probability distribution over downtime duration, d, which gives the probability of any duration of downtime being incurred. When considering downtime, negative downtime is equivalent to uptime and does not incur a cost in this problem formulation, and, therefore, $D_{c,i}$ is only of interest in the positive domain. The expected downtime of user agent i, $E(\text{downtime})$, given the probability distribution of downtime, $p(D_{c,i})$, is calculated by the integral on line 13. The solution to this integral when using Gaussian distributed random variables is presented in Section 5.4. The expected downtime is then added to the running total of weighed downtime on line 14.

The resource level of the user agent before the replenishment begins, $L_{u,i}$, is calculated on line 15, where the # operator denotes that the distribution has been adjusted to account for hard constraints on the state of the system. Note that this adjustment is only required when using probability distributions that do not already fit within the hard constraints (i.e. probability distributions that have infinite domain). A method for adjusting Gaussian distributions against hard constraints is discussed in Section 5.5. The quantity of the resource required to fully replenish the user agent, $Q_{u,i}$, is given on lines 16 and 17. Here, $* \leq L_a$ means that $Q_{u,i}$ is adjusted so that it does not exceed the resource level of the replenishment agent, L_a, as the replenishment agent may have insufficient supply of the resource to fully replenish

the user agent. A method for performing this soft adjustment for Gaussian distributions is outlined in Section 5.6.

The time taken to replenish the user agent is then given on line 18. The new levels of the user agent and replenishment agent after the replenishment are calculated on lines 19 and 20 respectively. Note that the unadjusted $Q_{u,i}$ is used on line 20 as using $Q^*_{u,i}$ can underestimate the amount of the resource transferred by the replenishment agent. The # adjustment ensures that L_a remains non-negative. The time that the user agent was last replenished, $T_{f,i}$, is updated on line 21, and the time that the replenishment agent finishes packing up, T_l, is given on line 22. The task is then removed from the schedule on line 23.

The final block of the algorithm (lines 24-28) calculates whether any of the user agents incur additional downtime between when they are last replenished and the completion time of the schedule. The algorithm concludes by returning the predicted resource levels of the agents (line 29) and the ratio objective function (line 30). The duration of the schedule, given by $T_l - t_{cur}$, is a random variable, and the expected value used on line (30) is simply the mean value of the probability distribution describing $T_l - t_{cur}$.

5. Using Gaussian distributed random variables

Analytical methods exist for calculating the sum of Gaussian distributed random variables, but do not exist for multiplication, division, and other operations used in the prediction framework presented in the previous section. To enable the use of Gaussian distributed random variables, this section introduces approximations to the inverse of a Gaussian distributed random variable, ratio of Gaussian distributed random variables, and product of Gaussian distributed random variables, as well as methods for evaluating the expected downtime, and adjusting Gaussian distributed random variables against hard and soft constraints.

5.1. Inverse Gaussian distributed random variable

A Gaussian approximation to the inverse of a Gaussian distributed random variable was presented in [7]. A new Gaussian approximation is introduced here which will be shown in Section 7 to outperform the one presented in [7]. Consider an inverse Gaussian distributed variable, I, that is formed by:

$$I = \frac{c}{G} \tag{4}$$

where G is a Gaussian distributed variable with mean u_G and standard deviation σ_G, and c is a constant. A Gaussian approximation of I can be attained by assuming that the points at $\mu_G + \sigma_G$ and $\mu_G - \sigma_G$ give the equivalent points at $\mu_I - \sigma_I$ and $\mu_I + \sigma_I$ respectively, where μ_I and σ_I are the mean and standard deviation of the Gaussian approximation of I. These give the following values for μ_I and σ_I:

$$\begin{aligned} \mu_I &= \frac{1}{2}\left(\frac{c}{\mu_G - \sigma_G} + \frac{c}{\mu_G + \sigma_G}\right) \\ &= \frac{c\mu_G}{\mu_G^2 - \sigma_G^2} \end{aligned} \tag{5}$$

$$\begin{aligned} \sigma_I &= \frac{1}{2}\left(\frac{c}{\mu_G - \sigma_G} - \frac{c}{\mu_G + \sigma_G}\right) \\ &= \frac{c\sigma_G}{\mu_G^2 - \sigma_G^2} \end{aligned} \tag{6}$$

When σ_G is small in comparison to μ_G, the inverse Gaussian distributed variable is highly Gaussian in shape. As σ_G is increased, the resultant inverse Gaussian distributed variable is skewed further to the right. The advantage of the approximation developed here over the one previously presented in [7] is demonstrated in Figure 2 for a case with high uncertainty. As can be seen, the new approximation better approximates the inverse Gaussian distribution, particularly in the left tail where the old approximation significantly overestimates the probability.

5.2. Ratio of Gaussian distributed variables

A method for approximating the ratio of two Gaussian distributed variables is given in [33]. A ratio, R:

$$R = \frac{E}{F} \tag{7}$$

with $E \sim \mathcal{N}(\mu_E, \sigma_E)$ and $F \sim \mathcal{N}(\mu_F, \sigma_F)$, and correlation between E and F of $\rho = 0$, can be approximated with a Gaussian distributed random variable where:

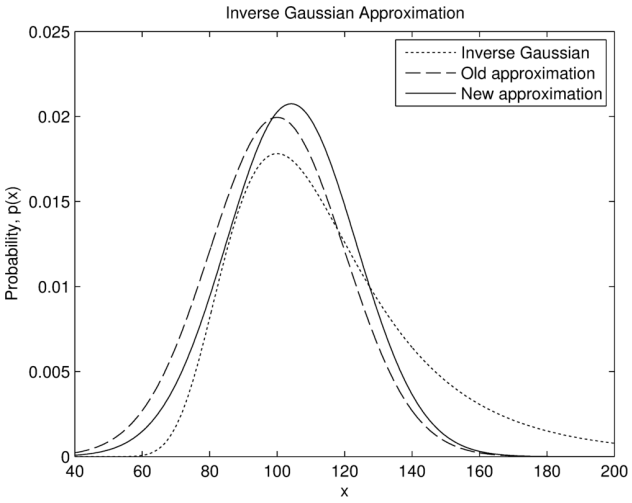

Figure 2: A comparison of the approximations of the inverse Gaussian distribution for $\mu_G/\sigma_G = 5$.

$$r = \frac{\sigma_F}{\sigma_E}, a = \frac{\mu_E}{\sigma_E} \text{ and } b = \frac{\mu_F}{\sigma_F} \tag{8}$$

$$\mu_R = \frac{a}{r(1.01b - 0.2713)} \tag{9}$$

$$\sigma_R = \frac{1}{r}\sqrt{\frac{a^2+1}{b^2+0.108b-3.795} - r^2\mu_R^2} \tag{10}$$

The authors specified that the approximation is only valid for $a < 2.5$, $b > 4$ [33]. As $a \to \infty$, the ratio of Gaussian distributed random variables is similar to an inverse Gaussian distributed random variable and can be approximated using the inverse Gaussian approximation presented above. For situations where $a \geq 2.5$, the inverse Gaussian approximation method with E treated as a scalar, $e = \mu_E$, has been used.

5.3. Product of Gaussian distributed variables

An approximation to the product of two Gaussian distributed variables is presented in [34]. For a product, M:

$$M = EF \tag{11}$$

Then:

$$\mu_M = \mu_E \mu_F \tag{12}$$

$$\sigma_M^2 = \sigma_E^2 \sigma_F^2 (1 + \delta_E^2 + \delta_F^2) \tag{13}$$

where

$$\delta_x = \frac{\mu_x}{\sigma_x} \tag{14}$$

The authors noted that the approximation improves as δ_E and δ_F become large [34].

5.4. Expected value

The integral required for calculating the risk-weighted downtime on line 12 of Algorithm 1 when using Gaussian distributions equates to:

$$\begin{aligned} E(D_{c,i}) &= \int_0^\infty d \; p(D_{c,i}) \; \mathrm{d}d \\ &= \int_0^\infty \frac{d}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(d-\mu)^2}{2\sigma^2}\right) \; \mathrm{d}d \\ &= \frac{\mu}{2}\left(1 + \operatorname{erf}\left(\frac{\mu}{\sigma\sqrt{2}}\right)\right) + \frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \end{aligned} \tag{15}$$

where μ and σ are the mean and standard deviation of the probability distribution describing the random variable $D_{c,i}$.

5.5. Adjusting against hard constraints

The probability distributions describing several random variables were required to be adjusted to take into consideration hard limitations on the system state; for example, the resource level of a user agent is bounded by 0 and $c_{u,i}$. A novel Gaussian approximation to the generalised rectified Gaussian distribution is introduced here. The rectified Gaussian distribution, used by [35], groups the probability in the negative domain at 0. The generalised rectified Gaussian distribution is proposed as an extension to this, where the distribution is rectified between two arbitrary values, a and b. If the original

CDF of the Gaussian distribution is $F(x)$, the problem is to calculate a new Gaussian PDF, $\mathcal{N}(\mu_R, \sigma_R^2)$, that approximates the PDF of the generalised rectified Gaussian distribution that satisfies the following CDF, $F_R(x)$:

$$F_R(x) = \begin{cases} 0 & \text{if } x < a \\ F(x) & \text{if } a \leq x < b \\ 1 & \text{if } x \geq b \end{cases} \tag{16}$$

where a and b are the limits on the state. The following process is similar to the truncation approach for constrained Kalman filtering [36]. The distribution being adjusted is first transformed to a standard normal distribution, yielding transformed constraints of c and d respectively:

$$c = \frac{a - \mu_A}{\sigma_A} \qquad d = \frac{b - \mu_A}{\sigma_A} \tag{17}$$

where μ_A and σ_A are the mean and standard deviation of the distribution being adjusted. The mean and variance of the Gaussian approximation of the rectified Gaussian distribution are then given by:

$$\begin{aligned} \mu_z &= \int_c^d \frac{\zeta}{\sqrt{2\pi}} \exp\left(-\frac{\zeta^2}{2}\right) \mathrm{d}\zeta + \frac{c}{2}\left(1 + \operatorname{erf}\left(\frac{c}{\sqrt{2}}\right)\right) \\ &\quad + \frac{d}{2}\left(1 - \operatorname{erf}\left(\frac{d}{\sqrt{2}}\right)\right) \\ &= \frac{1}{\sqrt{2\pi}}\left(\exp\left(\frac{-c^2}{2}\right) - \exp\left(\frac{-d^2}{2}\right)\right) \\ &\quad + \frac{c}{2}\left(1 + \operatorname{erf}\left(\frac{c}{\sqrt{2}}\right)\right) + \frac{d}{2}\left(1 - \operatorname{erf}\left(\frac{d}{\sqrt{2}}\right)\right) \end{aligned} \tag{18}$$

$$
\begin{aligned}
\sigma_z^2 &= \int_c^d (\zeta - \mu)^2 \exp\left(-\frac{\zeta^2}{2}\right) \mathrm{d}\zeta \\
&\quad + \frac{(c-\mu)^2}{2}\left(1 + \operatorname{erf}\left(\frac{c}{\sqrt{2}}\right)\right) \\
&\quad + \frac{(d-\mu)^2}{2}\left(1 - \operatorname{erf}\left(\frac{d}{\sqrt{2}}\right)\right) \\
&= \frac{\mu^2+1}{2}\left(\operatorname{erf}\left(\frac{d}{\sqrt{2}}\right) - \operatorname{erf}\left(\frac{c}{\sqrt{2}}\right)\right) \\
&- \frac{1}{\sqrt{2\pi}}\left(\exp\left(-\frac{d^2}{2}\right)(d-2\mu) - \exp\left(-\frac{c^2}{2}\right)(c-2\mu)\right) \\
&+ \frac{(c-\mu)^2}{2}\left(1 + \operatorname{erf}\left(\frac{c}{\sqrt{2}}\right)\right) \\
&+ \frac{(d-\mu)^2}{2}\left(1 - \operatorname{erf}\left(\frac{d}{\sqrt{2}}\right)\right)
\end{aligned} \tag{19}
$$

Taking the inverse of the transformation gives:

$$
\mu_R = \mu_z \sigma_A + \mu_A \qquad \sigma_R^2 = \sigma_z^2 \sigma_A^2 \tag{20}
$$

5.6. Adjusting against soft constraints

The other type of adjustment used adjusts one random variable so that it does not exceed another random variable. This soft adjustment is denoted by a * followed by the variable that it is adjusted against. This is used on line 17 of Algorithm 1 to ensure that the quantity of the resource used to replenish a user agent does not exceed the current capacity of the replenishment agent. Consider a random variable, A, that is adjusted so that it does not exceed the random variable B. The proposed method is as follows:

$$\mu_A^{*\leq B} = \begin{cases} \mu_A & \text{if } \mu_A - 3\sigma_A < \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A < \mu_B + 3\sigma_B \\ \mu_B & \text{if } \mu_A - 3\sigma_A > \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A > \mu_B + 3\sigma_B \\ \frac{\mu_B - 3\sigma_B + \mu_A + 3\sigma_A}{2} & \text{if } \mu_A - 3\sigma_A > \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A < \mu_B + 3\sigma_B \\ \frac{\mu_A - 3\sigma_A + \mu_B + 3\sigma_B}{2} & \text{if } \mu_A - 3\sigma_A < \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A > \mu_B + 3\sigma_B \end{cases} \tag{21}$$

$$\sigma_A^{*\leq B} = \begin{cases} \sigma_A & \text{if } \mu_A - 3\sigma_A < \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A < \mu_B + 3\sigma_B \\ \sigma_B & \text{if } \mu_A - 3\sigma_A > \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A > \mu_B + 3\sigma_B \\ \frac{\mu_A + 3\sigma_A - (\mu_B - 3\sigma_B)}{6} & \text{if } \mu_A - 3\sigma_A > \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A < \mu_B + 3\sigma_B \\ \frac{\mu_B + 3\sigma_B - (\mu_A - 3\sigma_A)}{6} & \text{if } \mu_A - 3\sigma_A < \mu_B - 3\sigma_B \\ & \text{and } \mu_A + 3\sigma_A > \mu_B + 3\sigma_B \end{cases} \tag{22}$$

This method ensures that $P(A \leq x) \leq P(B \leq x)$ for x within 3 standard deviations of the mean of both A and B.

6. Optimisation methods

This section presents the Apparent Tardiness Cost (ATC) heuristic in Section 6.1, a simulated annealing meta-heuristic in Section 6.2, and a branch and bound method in Section 6.3. All of these methods are used within an MPC-like framework—after each task is performed, the optimisation method is rerun to calculate a new task to be performed. This replanning enables unexpected changes to the state of the system to be considered by the optimisation. Before the optimisation method is run, the resource level of the replenishment agent is first checked to see whether it is above a threshold, $l_{a,thresh}$. If the resource level is below the threshold, then the replenishment agent is immediately sent back to the replenishment point to be replenished. A threshold of 5% of maximum capacity was found to give good results in the scenarios considered in this paper.

The ATC heuristic calculates what the next task of the replenishment agent should be, while the simulated annealing and branch and bound approaches both consider a schedule of tasks. The finite horizon used in this paper is the number of tasks in the schedule. This is to enable fair comparison between the simulated annealing and branch and bound approaches.

6.1. ATC heuristic

The ATC heuristic used by Kaplan and Rabadi in [10] for the aerial refuelling problem is a combination of the Weighted Shortest Processing Time first (WSPT) and Minimum Slack first (MS) rules. It calculates priorities for each task based on the following formula:

$$\pi_i = \frac{w_i}{d_{l,i}} \phi_i \tag{23}$$

where π_i is the priority of task i determined by the heuristic, $d_{l,i}$ is the total duration of the task, and ϕ_i is the marginal cost of delay. The task with the highest priority is selected as the next task to be performed. The marginal cost of delay used in [10] combined soft and hard deadlines with a ready time for each task. The SCAR scenarios under consideration only have a soft deadline, yielding a marginal cost of delay of:

$$\phi_i = \exp\left(-\frac{\max(0, t_{d,i} - t_{b,i})}{k\overline{t_b}}\right) \tag{24}$$

where $t_{b,i}$ is the time at which the replenishment agent begins replenishing the user agent, $\overline{t_b}$ is the average start time for all possible replenishment tasks for that replenishment agent, $t_{d,i}$ is the deadline for the user agent, and k is a scaling factor. The scaling factor biases the behaviour of the ATC heuristic towards the WSPT rule if k is very large, and towards the MS rule if k is very small. Typical values of k used range between 1 and 7. It should be noted that the deadline used by [10] is the time by which the task must be completed, whereas the deadline in a SCAR scenario is the time before which the replenishment agent must begin replenishing the user agent. The deadline is calculated for each user agent i as:

$$t_{d,i} = \frac{l_{u,i}}{r_{u,i}} \tag{25}$$

where $r_{u,i}$ is the mean value of $R_{u,i}$. The start time for replenishing each user agent i is calculated as:

$$t_{b,i} = \frac{s_{au,i}}{v_a} + d_{sa} \tag{26}$$

where v_a is the mean value of V_a, and d_{sa} is the mean value of D_{sa}. The total duration for each user agent, $d_{l,i}$, is given by:

$$d_{l,i} = t_{b,i} + \frac{c_{u,i} - \max\left(0,\ l_{u,i} - r_{u,i}t_{b,i}\right)}{r_a - r_{u,i}} + d_{pa} \tag{27}$$

where $r_{u,i}$ is the mean value of $R_{u,i}$, r_a is the mean value of R_a, and d_{pa} is the mean value of D_{pa}.

6.2. Simulated annealing

The simulated annealing method used by Kaplan and Rabadi [10, 11] was implemented, using the ATC heuristic to generate an initial schedule. The algorithm moves to neighbour solutions by randomly replacing a task in the schedule with one of the other possible tasks. The following limitation was placed on the generated schedule, $\boldsymbol{\theta}$:

$$\boldsymbol{\theta}_{i-1} \neq \boldsymbol{\theta}_i \neq \boldsymbol{\theta}_{i+1} \quad \forall i \in \{2, 3...a-1\} \tag{28}$$

where a is the number of tasks in the schedule. This ensures that successive tasks are different. The inputs to the simulated annealing algorithm are an initial temperature coefficient, a temperature cooling coefficient, a maximum number of inner loop iterations, and a maximum number of iterations. Values for these parameters are suggested in [10]. For the two scenarios examined in Section 7, a maximum number of inner loop iterations of 25, and a maximum number of iterations of 2000 worked well. For Scenario 1, an initial temperature coefficient of 0.2, and a temperature cooling coefficient of 0.95 gave good results, while in Scenario 2 an initial temperature coefficient of 0.1, and a temperature cooling coefficient of 0.7 performed well. The cost of each schedule is evaluated using the prediction framework from Section 4.

6.3. Branch and bound

The set of all possible schedules for a given finite horizon forms a tree where each branch represents the possible choices for the next task. Branch and bound minimises the size of the state space that is explored by culling branches of the tree where the minimum possible cost is higher than the cost of the best solution found so far [32]. Similar to the simulated annealing

implementation, the branch and bound implementation restricts consecutive tasks to be different. In addition, if the resource level of the replenishment agent is below the threshold $l_{a,thresh}$ at any point in the tree, the only valid task to be performed next is for the replenishment agent to be replenished by the replenishment point.

An important aspect of branch and bound is estimating the lower bound on the cost for each node. This lower bound represents the lowest possible cost of a complete schedule starting with the sequence of tasks described by that node. The more accurate the estimate of the lower bound, the more branches that can be pruned. A heuristic for estimating the minimum possible cost from any node in the tree for a SCAR scenario was developed in the authors' previous work [8] and has been used in this paper. Essentially, it assumes that the user agents do not exhaust their supply of the resource, while also providing a conservative estimate of the total time of the schedule. This guarantees that the estimated minimum cost of a schedule is below the actual cost.

To improve the search speed of the algorithm, the ATC heuristic was used to generate priorities for the tasks branching from each node. Two different methods of exploring the tree, shown in Figure 3, were considered. The first method, bottom-first, involves searching through the leaves first, and then gradually searching higher in the tree. This method has the advantage of not requiring any data to be stored in a tree structure, but has the disadvantage of focussing on one branch initially. The other method, top-first, explores the tree using a top down approach—the nodes are explored in priority order with changes initially occurring at the top level. As can be seen, each successive solution examined is in the opposite high level branch to the previous solution, ensuring that the breadth of the tree is explored rapidly. However, the top-first method requires the calculated costs and lower bounds of every node visited to be stored in a tree, which results in a memory complexity of $O(n^a)$, where n is the number of user agents and a is the number of tasks in the schedule.

The cost of the current best schedule versus the number of nodes explored for the two methods is compared in Figure 4 for a sample scenario. Both methods initially examine the same schedule generated by the ATC heuristic before searching other areas of the tree. Where the bottom-first approach finds neighbour schedules which make minor incremental improvements to the cost, the top-first approach quickly finds substantially better schedules in other branches of the tree. The bottom-first approach has many desir-

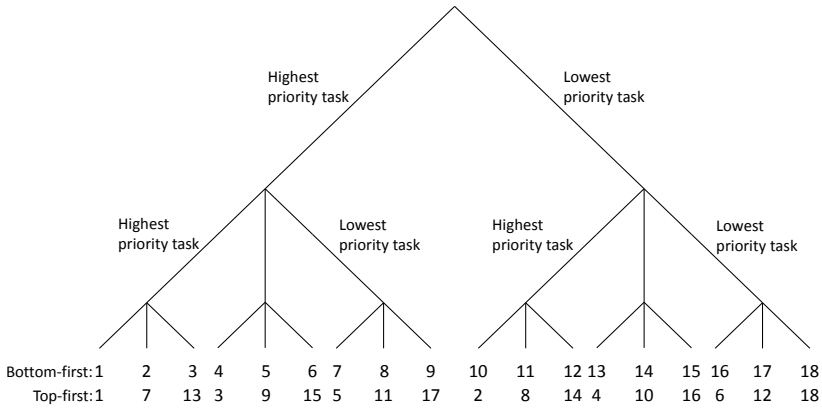
Highest priority task
Lowest priority task
Highest priority task
Lowest priority task
Highest priority task
Lowest priority task
Bottom-first: 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18
Top-first: 1 7 13 3 9 15 5 11 17 2 8 14 4 10 16 6 12 18

Figure 3: Two different methods for searching through a tree. The top line shows the exploration order using the bottom-first method, while the bottom line shows the exploration order using the top-first method.

able characteristics for small optimisation problems—low memory usage and minimal computational overhead associated with having to search through the tree. In larger problems, however, it may be computationally intractable to search through the entire tree and the anytime characteristic of branch and bound must be exploited. In these cases, the top-first approach is more desirable as it generally finds lower cost schedules for the same number of nodes explored as the bottom-first approach. In addition, since it focusses on earlier tasks, it fits quite well into the MPC-like framework used in this paper—optimisation efforts are focussed on the next tasks to be performed rather than the tasks at the end of the schedule. The top-first approach is used for the remainder of this paper.

To take advantage of the anytime characteristic of branch and bound, an optimisation depth was specified. As shown in Figure 5, the optimisation depth determines how far through the tree the branch and bound searches before selecting the remaining tasks using the ATC heuristic. This enables the computation time of the algorithm to be restricted while still using a long finite horizon. If the optimisation depth is equal to the schedule length, branch and bound will return the optimal schedule for that schedule length.

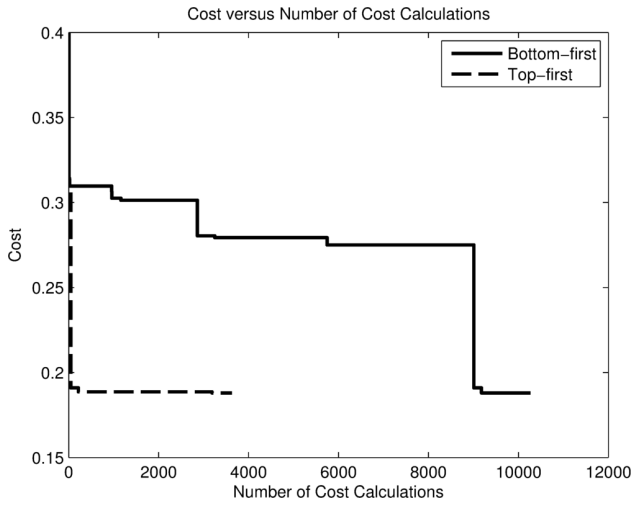

Figure 4: Cost versus number of calculations for the two branch and bound exploration methods. Lower costs are better. The bottom-first method incrementally improves on the solution, while the top-first method very quickly finds better solutions in other branches of the tree. In this example, the bottom-first method found the optimal schedule in 9178 calculations and required a total of 10269 calculations to fully explore the tree. The top-first method found the optimal schedule in 3186 calculations and required only 3629 calculations in total.

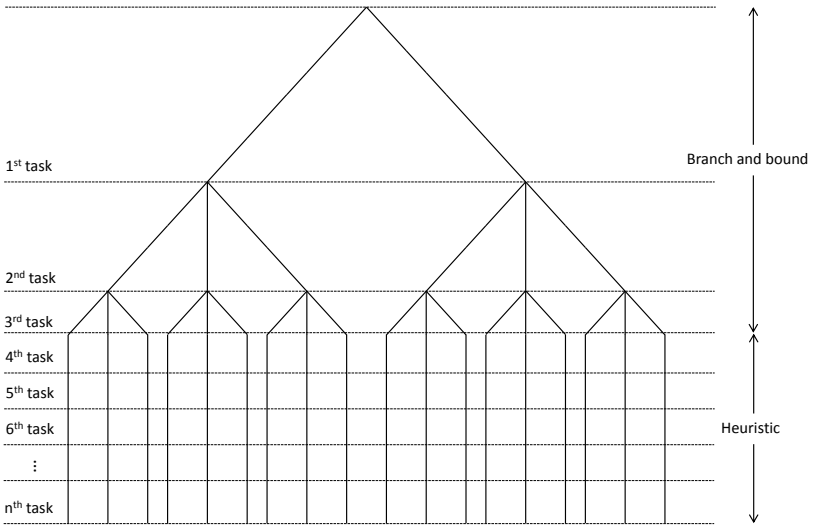

Figure 5: Branch and bound when the optimisation depth is smaller than the schedule length—the tasks within the optimisation depth (3 tasks in this example) are optimised using branch and bound, while the remaining tasks are selected using the ATC heuristic.

7. Computational study

This section evaluates the prediction and optimisation methods developed in this paper. All methods were implemented by the authors in Python on a 2.8GHz Intel i7-640M. Section 7.1 first introduces the two scenarios used to evaluate the methods. Section 7.2 then evaluates the prediction method developed in Section 4. Finally, Section 7.3 compares the developed branch and bound approach with the existing ATC heuristic and simulated annealing approaches.

7.1. Scenarios

7.1.1. Scenario 1

This scenario consists of several drills and excavators operating on specific benches within a mine site. The network of roads connecting the benches and replenishment point is shown in Figure 6. This road network is first reduced to a simpler graph without affecting the transit times by removing edges that do not form part of the shortest path between any two operational areas or the replenishment point. This simplified graph is shown in Figure

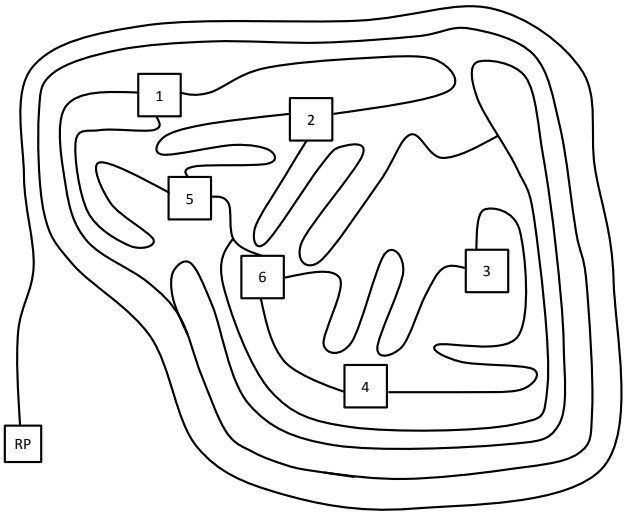

Figure 6: Example mine layout and road network used for Scenario 1. The locations of the Replenishment Point (RP) and the operational areas of the user agents (numbers correspond to the index of the user agent) are indicated by the squares.

7, with edge costs representing the distance between nodes. Note that the replenishment agent can pass through the operational area of a user agent without replenishing the user agent. The inset in Figure 7 shows the expected motion of a user agent. In this case, it is a drill that is drilling a specified hole pattern. While it is drilling a hole, it is stationary. Moving between holes is a very small proportion of the operating time of the drill.

The parameters of the user agents are shown in Table 1. The replenishment agent parameters are shown in Table 2, and the parameters of the replenishment point are shown in Table 3. The scenario was tested using the first 4, 5, and 6 user agents representing an under-utilised, fully-utilised, and over-utilised scenario respectively. In the under-utilised scenario the replenishment agent is operating below its capacity and should be able to prevent all of the user agents from exhausting their supply of the resource, while in the over-utilised scenario the replenishment agent is operating above its capacity. The fully-utilised case sits between the other two.

7.1.2. Scenario 2

The second scenario involves the delivery of fuel to 20 agents by truck. This number of agents is representative of large scenarios in the mining

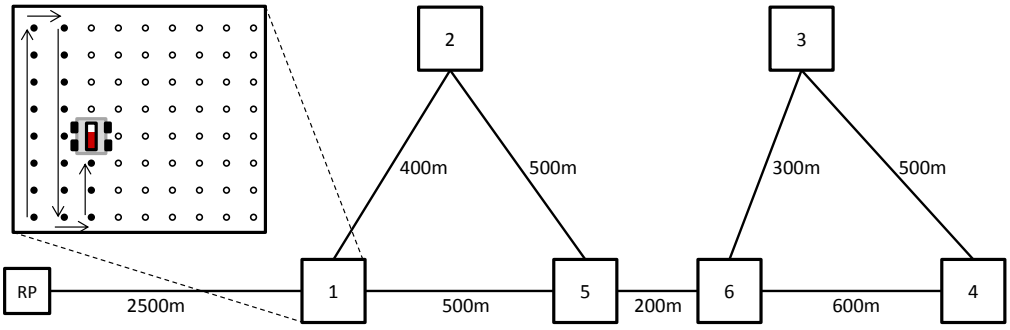

Figure 7: Locations of the Replenishment Point (RP), the operational areas of the user agents (numbers correspond to the index of the user agent), and the simplified graph for Scenario 1. Inset: Drill hole pattern on a bench. Solid circles represent holes that have been drilled, and empty circles are yet to be drilled. While drilling, the drill is stationary at the hole location.

Table 1: User Agent Parameters for Scenario 1

Agent	1	2	3	4	5	6
c_u (L)	1000	1200	700	1200	1000	800
R_u mean (L/s)	0.5	0.5	0.3	0.5	0.4	0.4
R_u standard deviation (L/s)	0.05	0.05	0.05	0.02	0.08	0.04

Table 2: Replenishment Agent Parameters for Scenario 1

Parameter	Mean	Standard Deviation
c_a (L)	5000	–
R_a (L/s)	10	0.5
D_{sa} (s)	60	20
D_{pa} (s)	20	5
V_a (m/s)	15	0.5

Table 3: Replenishment Point Parameters for Scenario 1

Parameter	Mean	Standard Deviation
D_{sr} (s)	30	10
D_{pr} (s)	10	1
R_r (L/s)	20	1

Table 4: Replenishment Agent Parameters for Scenario 2

Parameter	Mean	Standard Deviation
c_a Large (L)	2760	–
c_a Medium (L)	1800	–
c_a Small (L)	1200	–
R_a (L/hr)	720	72
D_{sa} (min)	12	1
D_{pa} (min)	4	0.3
V_a (km/hr)	16	4

and agricultural domains. The quality of roads between operational areas is highly variable which results in travel speeds that are very uncertain. The agents use the fuel at a relatively predictable rate in comparison to the uncertainty of the speed of the truck. Figure 8 shows the operational areas of the user agents and the simplified graph of the roads connecting them. The full road network has been omitted in the interest of space.

Three different size replenishment agents were tested, and their parameters are shown in Table 4. The various sizes of the replenishment agent only vary in their capacity—the resource usage rates, set-up and pack-up times, and velocity, are the same for all sizes. The large, medium, and small replenishment agent sizes roughly correspond to under-, fully-, and over-utilised scenarios respectively. The parameters of the replenishment point are shown in Table 5. There are four different sizes of user agents, each with 2 days supply of fuel. The parameters of each size of user agent are shown in Table 6, and the size of each user agent is outlined in Table 7.

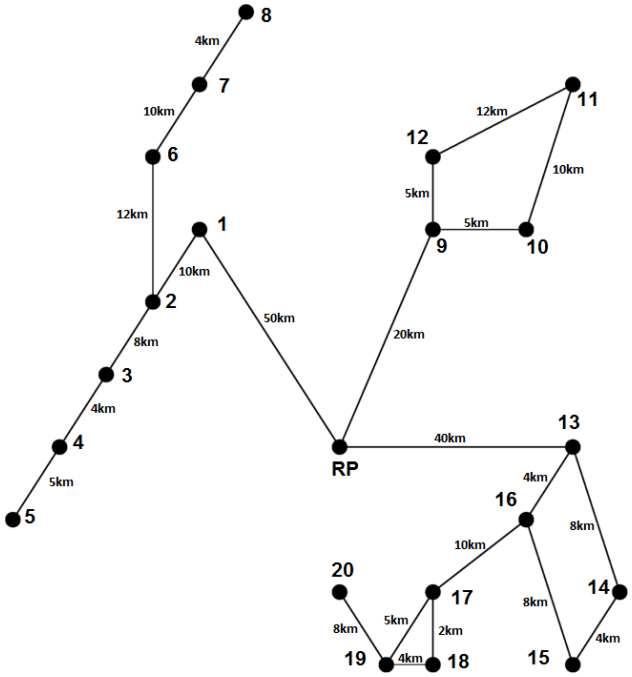

Figure 8: Layout of the 20 user agents in Scenario 2. Road distances are shown in km.

Table 5: Replenishment Point Parameters for Scenario 2

Parameter	Mean	Standard Deviation
D_{sr} (min)	12	2
D_{pr} (min)	6	1
R_r (L/hr)	12000	1200

Table 6: User Agent Parameters for Scenario 2

Type	Parameter	Mean	Standard Deviation
Small (S)	c_u (L)	480	–
	R_u (L/hr)	10	1
Medium (M)	c_u (L)	600	–
	R_u (L/hr)	12.5	1.25
Large (L)	c_u (L)	720	–
	R_u (L/hr)	15	1.5
Extra Large (XL)	c_u (L)	960	–
	R_u (L/hr)	20	2

Table 7: User Agent Types for Scenario 2

Agent	1	2	3	4	5	6	7	8	9	10
Size	S	L	S	M	M	S	M	L	XL	L
Agent	11	12	13	14	15	16	17	18	19	20
Size	L	XL	M	L	M	S	L	M	L	M

7.2. Evaluation of the prediction framework

The prediction method developed in Section 4 was compared with the analytical prediction method previously developed by the authors in [7], using a Monte Carlo generated cost as a benchmark. They were tested in Scenario 1 with 6 user agents and a schedule length of 8 tasks, and in Scenario 2 with the large replenishment agent and a schedule length of 20 tasks. 10,000 random schedules were generated for each scenario with the initial resource levels of all agents randomly initialised to a value between 0% and 100% of capacity for each schedule. The Monte Carlo cost was generated using 1,000 samples as this was found to be a good trade-off between error and calculation time in [7].

Table 8 shows the error of the proposed and previous methods in comparison to the Monte Carlo method. As can be seen, the proposed method has significantly less error than the previous method, particularly in Scenario 2. A significant source of the error in Scenario 2 for the previous method was from the particular approximation of the inverse Gaussian distributed random variable that was used. As the standard deviation of the velocity of the replenishment agent is very high compared to the mean, the resultant inverse Gaussian distribution for the travel time is heavily skewed. As was shown in Figure 2, the old approximation of the inverse Gaussian distribution overestimates the probability in the left tail of the distribution in these cases, resulting in the cost being underestimated. The new approximation does not overestimate this probability to the same extent, and consequently produces significantly better results.

The comparison accuracy shows how effective each method is at discriminating between schedules. This is the most important aspect of the prediction method—it must be able to accurately discriminate between schedules for it to be effective when used within a schedule optimisation. The proposed method is accurate in 0.3% more cases in Scenario 1, and 1% more cases in Scenario 2. In Scenario 2 in particular, this is a significant improvement.

The main advantage of the proposed prediction method over the Monte Carlo method is the computation time. The proposed method took just 7ms in Scenario 1 compared to 678ms for the Monte Carlo method, and in Scenario 2 took 9ms compared to 1.506s for the Monte Carlo method.

7.3. Evaluation of the optimisation methods

The optimisation methods tested are summarised below, and the acronyms in parentheses will be used to refer to the methods.

Table 8: Proposed cost method minus Monte Carlo cost method

Scenario	Method	Mean ($\times 10^{-3}$)	Standard Deviation ($\times 10^{-3}$)	Comparison Accuracy
1	Previous from [7]	3.18	3.11	99.3%
1	Proposed	0.08	1.52	99.6%
2	Previous from [7]	-27.8	4.94	98.4%
2	Proposed	-1.96	1.92	99.4%

- ATC heuristic (ATC)
- Simulated annealing using the developed prediction framework (SA)
- Branch and bound ignoring uncertainty (DBB)
- Branch and bound incorporating uncertainty through the developed prediction framework (SBB)

DBB uses the Monte Carlo method from [7] with 1 sample, treating all parameters as certain. The SA and SBB methods incorporate uncertainty through the developed prediction framework.

7.3.1. Scenario 1

Each simulation of Scenario 1 was initialised with random initial resource levels between 50% and 100% to simulate realistic in-progress starting conditions, and the simulation lasted for 5 hours of simulated time. The k value for ATC was first tuned by running multiple simulations with values between 1 and 7. As shown in Figure 9, the lowest costs were achieved using a k value of approximately 2.5 for the 4-user agent scenario, and approximately 5.5 for the 5- and 6-user agent scenarios. This means that the behaviour is biased more towards the MS rule than the WSPT rule for the 4-user agent scenario in comparison to the 5- and 6-agent scenarios.

Each optimisation method was tested 40 times to account for the variability due to the stochastic nature of the simulation. The SA, DBB, and SBB methods were tested using a schedule length of $n + 3$ tasks, where n is the number of user agents in the system. Using shorter schedule lengths than this can lead to undesirable behaviour; for further discussion, see [8].

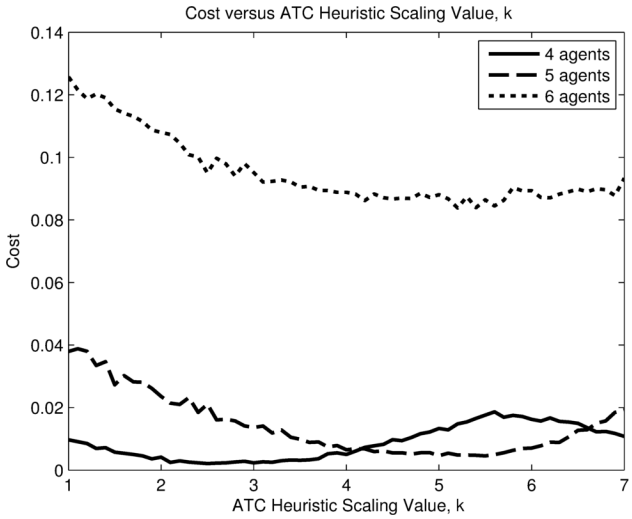

Figure 9: Cost versus k value for the ATC heuristics for the 4-, 5-, and 6-user agent cases in scenario 1.

The percentage downtime results from this scenario are shown in Figure 10, and Figure 11 shows the percentage of simulations in which none of the user agents exhausted their supply of the resource. As can be seen, the proposed SBB method consistently produced the lowest downtime in the 4- and 5-user agent scenarios. In the 6-user agent scenario, the SBB and DBB methods produced almost identical performance. In over-utilised scenarios like this, the distributions for the downtime of the user agents are predominantly in the positive domain. The expected cost of these distributions is therefore very close to the mean value, and hence very similar to the result returned by the deterministic framework used in DBB. The SA method struggled to find good schedules, providing only a minimal improvement to the initial schedule generated by the ATC heuristic.

The main advantage of the SBB method over DBB is highlighted by the percentage of simulations in which none of the user agents incurred downtime. In the 4- and 5-user agent scenarios, the schedules tested by DBB will frequently have a cost of zero. This means that, in many cases, DBB is unable to differentiate between these schedules and consequently relies on the priorities generated by the ATC heuristic to select a good schedule. The proposed prediction framework used by SBB will never result in a cost of zero as there is always some risk associated with each schedule. This is illustrated

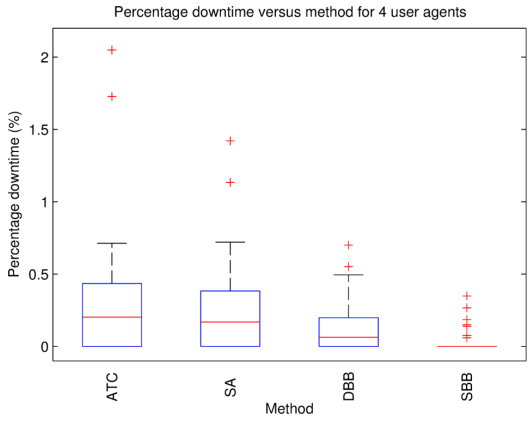

(a) 4 user agents

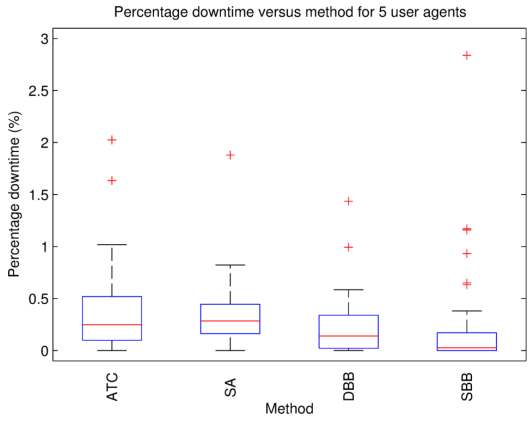

(b) 5 user agents

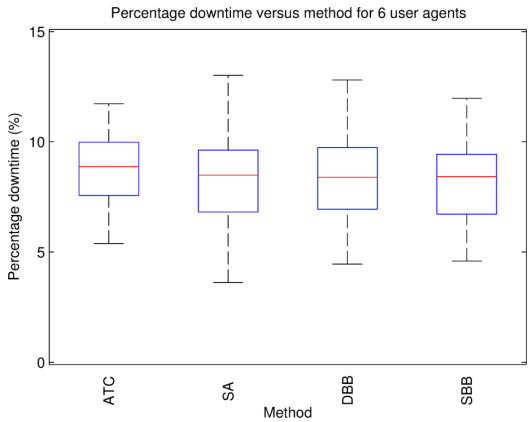

(c) 6 user agents

Figure 10: Box and whisker plots for the percentage downtime in Scenario 1. Lower results are better.

Table 9: Calculation times in seconds for Scenario 1

Number of agents, n	ATC	SA	DBB	SBB
4	1.37E-4	6.55	2.80E-2	0.884
5	1.48E-4	7.96	4.60E-2	2.51
6	1.84E-4	9.68	7.46	37.8

in Figure 12. SBB will find that selecting a schedule that will replenish the user agent at point a is significantly less risky than at point b, whereas DBB will return a cost of zero for both schedules and is unable to differentiate between them. Therefore, point b may be chosen sometimes by DBB, leading to an incurred cost when the actual resource level is as shown in Figure 12.

The calculation times for the various optimisation methods are detailed in Table 9. The ATC heuristic is the fastest of the methods, taking a fraction of a second in all cases, while SA consistently takes several seconds. DBB computes a solution very quickly in the 4- and 5-user agent scenarios because it can find a zero-cost schedule—it will usually find a zero-cost schedule in the first few schedules examined and will return this immediately. SBB takes longer in these cases as it searches through the entire tree. In the 6-user agent scenario, DBB is unable to find a zero-cost schedule and spends significantly longer searching through the tree. In this case, SBB is approximately 5 times slower than DBB. This increase reflects the increased computational requirement of the proposed prediction framework over frameworks that ignore uncertainty.

7.3.2. Scenario 2

Each optimisation method was tested 40 times in this scenario, with initial conditions between 50% and 100% randomly selected and each simulation lasting for 9 days of simulated time. ATC scaling values, k, of 3 for the large and medium sized replenishment agents, and 5 for the small replenishment agent, were found to give good behaviour. Given the large size of this scenario, calculating an optimal schedule is not feasible as there are over 10^{32} combinations when a schedule of 25 tasks is considered. This is where the anytime behaviour of branch and bound is a significant benefit over the A* method used in [8]. This anytime nature was exploited in two ways—the optimisation depth of the algorithm was varied, and a hard limit of 10,000 nodes was placed on the size of the solution tree.

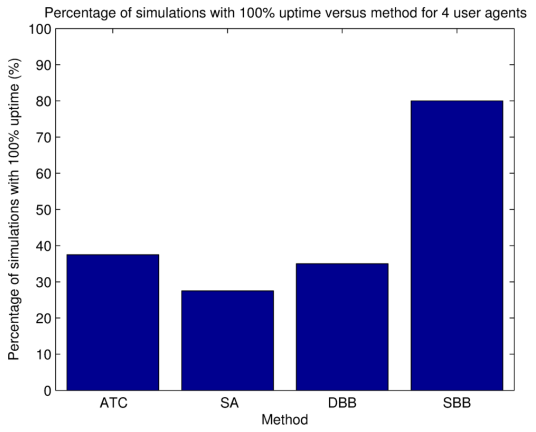

(a) 4 user agents

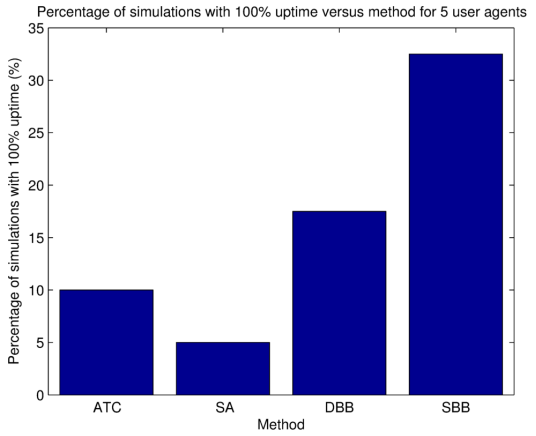

(b) 5 user agents

Figure 11: Percentage of simulation runs with 100% uptime in Scenario 1. Higher results are better. The 6-user agent scenario results are omitted as no method was able to achieve 100% uptime in any of the simulations.

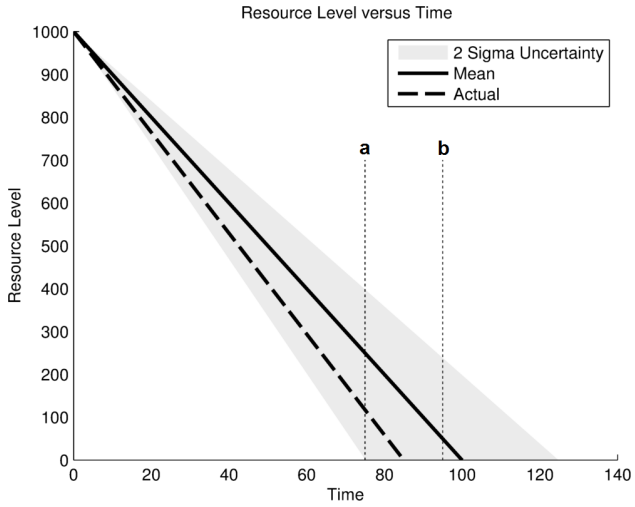

Figure 12: Level of a user agent showing the predicted resource level, an uncertainty of two standard deviations (2 sigma), and the actual resource level. Replenishing at point a is selected when considering uncertainty, while the two points cannot be differentiated if uncertainty is not considered.

Selecting the schedule length is an important aspect of this problem. Using a short schedule length can lead to myopic behaviour, while using a long schedule length exponentially increases the size of the solution tree. When using an optimisation depth that is smaller than the schedule length, the remaining tasks in the schedule are selected by the ATC heuristic. If the schedule length is too long in comparison to the optimisation depth, the suboptimal choices of the heuristic can also reduce the benefit of using a longer schedule length. Figure 13 shows the percentage of simulations with 100% uptime using the SBB algorithm with an optimisation depth of two tasks as a function of the schedule length. As can be seen, the best performance is achieved at a schedule length of 25 tasks.

Figures 14 and 15 show the results for the optimisation methods in the second scenario. SA, DBB, and SBB were tested using a schedule length of 25 tasks, with optimisation depths of 1, 2, and 3 used by DBB and SBB. These results broadly mirror those in Scenario 1, with the proposed SBB method producing the best results. The benefit of the directed optimisation of the branch and bound methods on an initial schedule generated by the ATC heuristic is evident here. Even if only the first task is optimised, DBB and SBB provide a huge benefit over the ATC heuristic. In both the large and

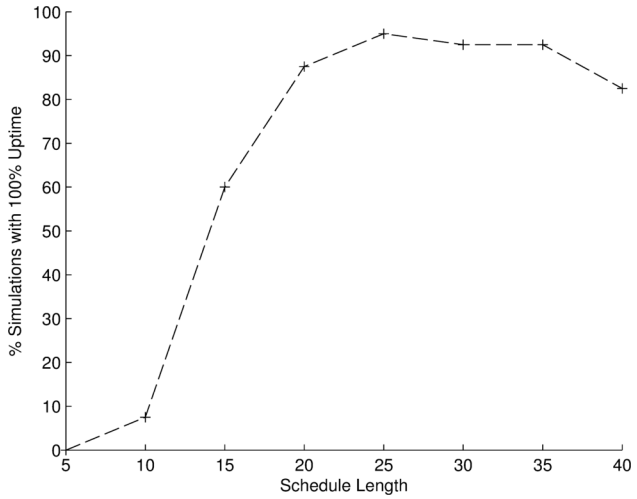

Figure 13: Percentage of simulations with 100% uptime for various schedule lengths using SBB optimising the first two tasks in scenario 2. Higher results are better.

medium replenishment agent cases, SBB clearly outperformed DBB. This is highlighted by the results in Figure 15. In the small replenishment agent case, SBB and DBB produced similar results, corroborating the findings from Scenario 1. As these methods are used within an MPC-like framework, it is beneficial to focus the optimisation on earlier tasks within a schedule. SA struggled to find good schedules in this scenario as it spread the optimisation efforts across the entire schedule. As a result, it was unable to sufficiently explore the search space to yield much improvement over the ATC heuristic.

Larger scenarios than this can be handled by the proposed approach, but may require additional planning time to achieve the significant improvements seen in the two scenarios examined in this paper. This is because the number of possible schedules increases exponentially as the number of user agents is increased. If the search time is held constant, then the solution qualities from the branch and bound methods will gradually decrease with the number of user agents. In the limit, the behaviour of the algorithm will become dominated by the heuristic used to select tasks for generating complete schedules out to the planning horizon.

The calculation times for each method are shown in Table 10. These results follow similar trends to the results in Scenario 1, with the ATC heuristic computing very quickly, and the branch and bound methods taking the

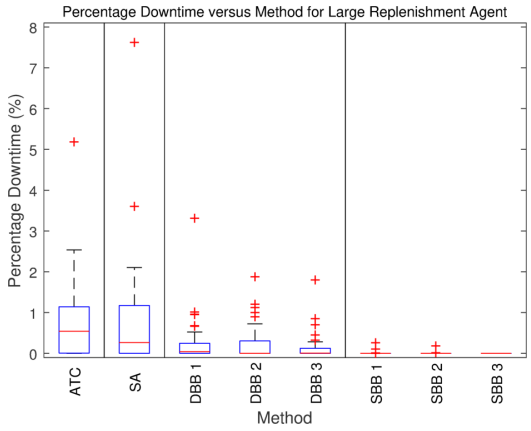

(a) Large replenishment agent

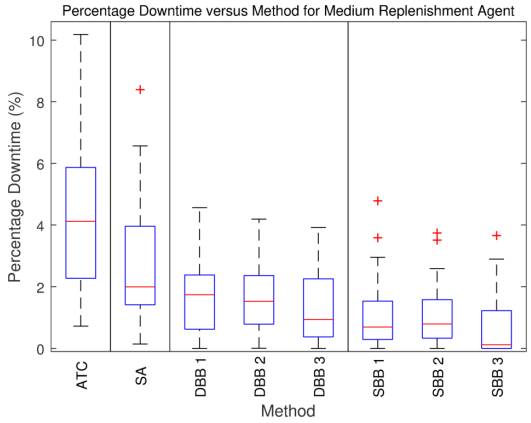

(b) Medium replenishment agent

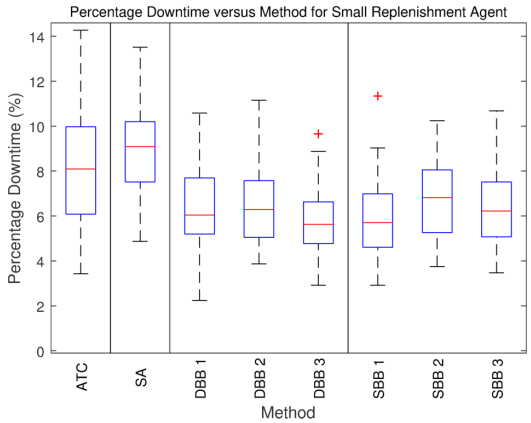

(c) Small replenishment agent

Figure 14: Box and whisker plots for the percentage downtime in Scenario 2. Lower results are better. The number next to the DBB and SBB methods is the optimisation depth.

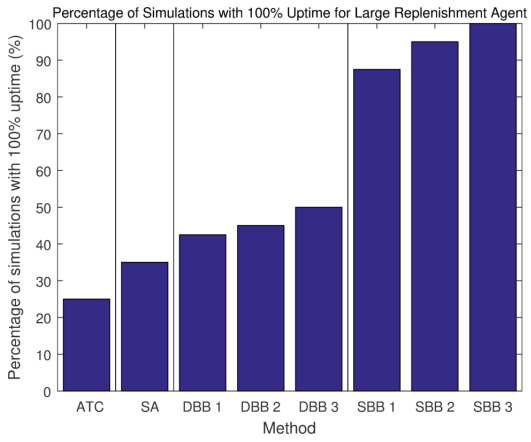

(a) Large replenishment agent

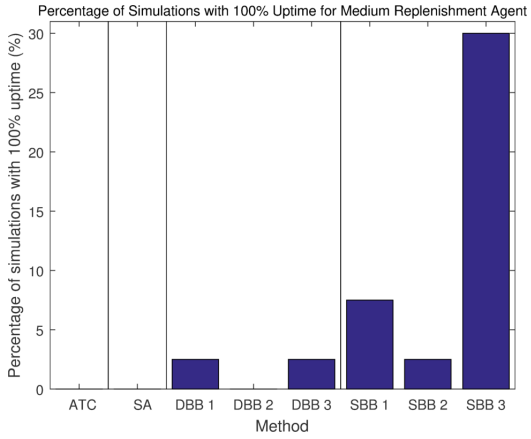

(b) Medium replenishment agent

Figure 15: Percentage of simulation runs with 100% uptime in Scenario 2. Higher results are better. The number next to the DBB and SBB methods is optimisation depth. The small replenishment agent scenario results are omitted as no method was able to achieve 100% uptime in any of the simulations.

Table 10: Calculation times in seconds for Scenario 2

Agent size	ATC	SA	DBB 1	DBB 2	DBB 3
Large	1.61E-3	13.5	0.664	0.674	0.668
Medium	1.50E-3	13.9	0.671	0.639	0.638
Small	1.64E-3	13.1	2.38	24.2	29.8
Agent size	SBB 1	SBB 2	SBB 3		
Large	1.38	11.9	12.2		
Medium	1.88	21.7	39.5		
Small	1.68	20.1	48.6		

longest. In the large and medium replenishment agent cases, DBB has a short calculation time compared to SBB as it quickly finds zero-cost schedules. In the small replenishment agent case, the calculation times of DBB and SBB are much closer together. Many opportunities exist for improving the calculation times of these methods including using a language such as C, using a parallel implementation of branch and bound, and storing more data in the tree structure to reduce the computation required at each node. Using conservative estimates, speed increases of at least 100 times are feasible, giving potential sub-second calculation times for the branch and bound methods.

8. Conclusion

Research on replenishment and collection scenarios has thus far not taken into account long term considerations, such as the limited capacity of the replenishment agent and replenishing user agents multiple times, that are critical for achieving persistent autonomy. More importantly, the examined literature has treated these scenarios as deterministic, ignoring the uncertainty that is inherent in realistic scenarios. These aspects play an important role in the optimisation process, and the SCAR scenario was developed to specifically address these shortcomings.

This paper proposed a novel framework for incorporating uncertainty when predicting the outcome of the schedule for the replenishment agent in a SCAR scenario, and developed a branch and bound method that used the prediction framework to optimise the schedule of the replenishment agent.

Improved Gaussian approximations enabled the proposed prediction framework to outperform an existing framework. The branch and bound approach using this framework was then shown to outperform the ATC heuristic, a simulated annealing meta-heuristic, and a branch and bound approach ignoring uncertainty, in both a small and large scenario. In the large scenario, the anytime characteristic of branch and bound was exploited to find good schedules within a reasonable length of time by varying the optimisation depth of the algorithm. Tasks beyond the optimisation depth were selected using the ATC heuristic, enabling the use of a long schedule length to reduce myopic decision making.

An interesting avenue of future work is considering systems with multiple replenishment agents or multiple resources, as the size of the search space will be substantially larger than for the single replenishment agent or single resource scenarios, thus requiring more efficient optimisation methods in order to select appropriate tasks. One possible method could be to cluster the user agents so that the problem reduces to multiple single-replenishment agent optimisations. Other areas of future work include considering uncertainty on the current state of the system, assessing the robustness of the methods to changes in the underlying probability distributions, and developing methods for dynamically adapting the k values for the ATC heuristic.

Acknowledgements

This work was supported by the Rio Tinto Centre for Mine Automation and the Australian Centre for Field Robotics, University of Sydney, Australia.

References

[1] J. Nguyen, N. Lawrance, R. Fitch, S. Sukkarieh, Energy-Constrained Motion Planning for Information Gathering with Autonomous Aerial Soaring, in: IEEE ICRA, 2013, pp. 3810–3816.

[2] D. Austin, L. Fletcher, A. Zelinsky, Mobile robotics in the long term-exploring the fourth dimension, in: IEEE/RSJ IROS, Vol. 2, 2001, pp. 613–618 vol.2. doi:10.1109/IROS.2001.976237.

[3] R. Luo, C. Liao, K. Su, K. Lin, Automatic docking and recharging system for autonomous security robot, in: Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on, 2005.

[4] J. Forlizzi, C. DiSalvo, Service robots in the domestic environment: a study of the roomba vacuum in the home, in: Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction, HRI '06, ACM, New York, NY, USA, 2006, pp. 258–265. doi:10.1145/1121241.1121286.

[5] InTouch Health, RP-7i Robot. http://www.intouchhealth.com/products-and-services/products/rp-7i-robot/ (2012).
URL http://www.intouchhealth.com/products-and-services/products/rp-7i-robot/

[6] P. Zebrowski, R. T. Vaughan, Recharging robot teams: A tanker approach, in: Proceedings of the 12th International Conference on Advanced Robotics, ICAR '05, 2005, pp. 803–810. doi:10.1109/ICAR.2005.1507500.

[7] A. W. Palmer, A. J. Hill, S. J. Scheding, Stochastic collection and replenishment (SCAR): Objective functions, in: Proceedings of the 2013 IEEE International Conference on Intelligent Robots and Systems (IROS), 2013, pp. 3324–3331. doi:10.1109/IROS.2013.6696829.

[8] A. W. Palmer, A. J. Hill, S. J. Scheding, Stochastic Collection and Replenishment (SCAR) Optimisation for Persistent Autonomy, in: Proceedings of the 2014 IEEE International Conference on Intelligent Robots and Systems (IROS), 2014, pp. 2943–2949. doi:10.1109/IROS.2014.6942968.

[9] Z. Jin, T. Shima, C. J. Schumacher, Optimal scheduling for refueling multiple autonomous aerial vehicles, IEEE Transactions on Robotics 22 (4) (2006) 682–693. doi:10.1109/TRO.2006.878793.

[10] S. Kaplan, G. Rabadi, Exact and heuristic algorithms for the aerial refueling parallel machine scheduling problem with due date-to-deadline window and ready times, Computers & Industrial Engineering 62 (1) (2012) 276–285. doi:10.1016/j.cie.2011.09.015.

[11] S. Kaplan, G. Rabadi, Simulated annealing and metaheuristic for randomized priority search algorithms for the aerial refuelling parallel machine scheduling problem with due date-to-deadline windows and

release times, Engineering Optimization 45 (1) (2013) 67–87. doi:10.1080/0305215X.2012.658783.

[12] J. W. Barnes, V. D. Wiley, J. T. Moore, D. M. Ryer, Solving the aerial fleet refueling problem using group theoretic tabu search, Mathematical and Computer Modelling 39 (6) (2004) 617–640. doi:10.1016/S0895-7177(04)90544-4.

[13] H. Shen, P. Tsiotras, Optimal Scheduling for Servicing Multiple Satellites in a Circular Constellation, in: Proceedings of the 2002 AIAA/AAS Astrodynamics Specialists Conference, 2002, pp. 1–10.

[14] B. Kannan, V. Marmol, J. Bourne, M. B. Dias, The Autonomous Recharging Problem: Formulation and a Market-based Solution, in: Proceedings of the 2013 IEEE International Conference on Robotics and Automation (ICRA), 2013, pp. 3488–3495.

[15] Y. Litus, P. Zebrowski, R. T. Vaughan, A distributed heuristic for energy-efficient multirobot multiplace rendezvous, IEEE Transactions on Robotics 25 (1) (2009) 130–135. doi:10.1109/TRO.2008.2007459.

[16] N. Mathew, S. L. Smith, S. L. Waslander, Multirobot Rendezvous Planning for Recharging in Persistent Tasks, IEEE Transactions on Robotics 31 (1) (2015) 128–142.

[17] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, P. Corke, Data collection, storage, and retrieval with an underwater sensor network, in: Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems - SenSys '05, 2005, pp. 154–165. doi:10.1145/1098918.1098936.

[18] Y. Tirta, Z. Li, Y.-H. Lu, S. Bagchi, Efficient collection of sensor data in remote fields using mobile collectors, in: Proceedings on the 13th IEEE International Conference on Computer Communications and Networks (ICCCN), 2004, pp. 515–519.

[19] B. Yuan, M. Orlowska, S. Sadiq, On the Optimal Robot Routing Problem in Wireless Sensor Networks, IEEE Transactions on Knowledge and Data Engineering 19 (9) (2007) 1252–1261. doi:10.1109/TKDE.2007.1062.

[20] O. Tekdas, V. Isler, Using mobile robots to harvest data from sensor fields, IEEE Wireless Communications 16 (1) (2009) 22–28.

[21] O. Tekdas, D. Bhadauria, V. Isler, Efficient data collection from wireless nodes under the two-ring communication model, The International Journal of Robotics Research 31 (6) (2012) 774–784. doi:10.1177/0278364912439429.

[22] D. Bhadauria, O. Tekdas, V. Isler, Robotic data mules for collecting data over sparse sensor fields, Journal of Field Robotics 28 (3) (2011) 388–404. doi:10.1002/rob.

[23] M. Ma, Y. Yang, M. Zhao, Tour planning for mobile data-gathering mechanisms in wireless sensor networks, IEEE Transactions on Vehicular Technology 62 (4) (2013) 1472–1483. doi:10.1109/TVT.2012.2229309.

[24] Y. Yan, Y. Mostofi, An Efficient Clustering and Path Planning Strategy for Data Collection in Sensor Networks Based on Space-Filling Curves, in: Proceedings of the 2014 IEEE Conference on Decision and Control, 2014, pp. 6895–6901.

[25] A. Allahverdi, C. Ng, T. Cheng, M. Y. Kovalyov, A survey of scheduling problems with setup times or costs, European Journal of Operational Research 187 (3) (2008) 985–1032. doi:10.1016/j.ejor.2006.06.060.

[26] C. Fang, P. Yu, B. C. Williams, Chance-Constrained Probabilistic Simple Temporal Problems, Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (2014) 2264–2270.

[27] E. Raboin, P. Švec, D. S. Nau, S. K. Gupta, Model-predictive asset guarding by team of autonomous surface vehicles in environment with civilian boats, Autonomous Robots 38 (2014) 261–282. doi:10.1007/s10514-014-9409-9.

[28] D. Bertsimas, M. Sim, Robust discrete optimization and network flows, Mathematical Programming 98 (1-3) (2003) 49–71. doi:10.1007/s10107-003-0396-4.

[29] S. Yoon, A. Fern, R. Givan, FF-Replan: A baseline for probabilistic planning, Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS) (2007) 352–359.

[30] J.-P. Arnaout, G. Rabadi, J. H. Mun, A dynamic heuristic for the stochastic unrelated parallel machine scheduling problem, International Journal of Operations Research 3 (2) (2006) 136–143.

[31] D. Ferguson, M. Likhachev, A. Stentz, A Guide to Heuristic-based Path Planning, in: Proceedings of the International Workshop on Planning Under Uncertainty for Autonomous Systems, International Conference on Automated Planning and Scheduling (ICAPS), 2005, pp. 9–18.

[32] A. Land, A. Doig, An automatic method of solving discrete programming problems, Econometrica: Journal of the Econometric Society 28 (3) (1960) 497–520.

[33] G. Marsaglia, Ratios of normal variables, Journal of Statistical Software 16 (4) (2006) 1–10.

[34] R. Ware, F. Lad, Approximating the distribution for sums of products of normal variables, Tech. Rep. 1978, The University of Queensland (2003).

[35] J. Meng, J. M. Zhang, Y. Chen, Y. Huang, Bayesian non-negative factor analysis for reconstructing transcription factor mediated regulatory networks, Proteome Science 9 (Suppl 1) (2011) S9. doi:10.1186/1477-5956-9-S1-S9.

[36] A. W. Palmer, A. J. Hill, S. J. Scheding, Applying Gaussian distributed constraints to Gaussian distributed variables, Information Fusiondoi:10.1016/j.inffus.2016.02.008.